

Calif. Bill Goes Too Far In Trying To Regulate Attorney AI Use

By **Joshua Wurtzel** (September 23, 2026)

Editor's note: Law360 welcomes opinionated commentary and debate in our Expert Analysis section. To submit op-eds or rebuttals, or to speak to an editor about submissions, please email expertanalysis@law360.com.

California S.B. 574 is the first state statute in the country governing how attorneys and arbitrators use generative artificial intelligence.

It bars attorneys from entering confidential information into AI systems unless access is restricted to the attorney and those the attorney authorizes, requires attorneys to disclose their use of AI to the court for all documents submitted to it, and prohibits delegating the practice of law to AI. A parallel provision imposes nondelegation and disclosure duties on arbitrators.



Joshua Wurtzel

The bill cleared the California Legislature on Aug. 31, without a single dissenting vote, and was presented to Gov. Gavin Newsom on Sept. 9. He has until Sept. 30 to act — though short of an affirmative veto, the bill will become law.

The bill's Senate Judiciary Committee analysis states that it "protects those receiving legal services by codifying certain safeguards for the use of AI by attorneys and arbitrators" — a response to the wave of sanctions against lawyers who filed briefs containing AI-hallucinated case citations. But this bill is, at best, useless, and at worst, detrimental to the practice of law.

Read narrowly, the bill restates duties lawyers already have. The Model Rules of Professional Conduct already require a lawyer to be competent, which includes confirming that authority in a brief exists and says what the brief claims. And they also already require a lawyer to protect client confidences, which includes knowing who can access client information. The bill unnecessarily layers on top of both.

But read broadly, the bill may ban lawyers and arbitrators from using AI at all — even fully supervised, fully edited AI use. And it also treats AI as inherently suspect. This unduly stigmatizes, and will potentially chill, AI usage that would otherwise improve both a lawyer's work quality and efficiency.

In other words, the provisions worth having arguably add nothing. The provisions that add something go too far.

The confidentiality provision doesn't fit how AI actually gets deployed.

The bill's new Business and Professions Code Section 6068.1(a)(3)(A) requires that an attorney using AI "[n]ot enter confidential, personal identifying, and other nonpublic information into" an AI system that does not restrict this information "to the attorney and persons authorized by the attorney" who also owe a duty to protect it.

That language turns on who can technically reach the data, not on how well the AI is protected. And for almost all hosted AI platforms, the answer is the same: Enterprise AI contracts protect client data through vendor-side confidentiality commitments and data

processing agreements.

The people with access to a prompt once it reaches Anthropic, OpenAI, Harvey, Microsoft or Legora aren't people the attorney authorized. They're the vendor's engineers, support staff and subprocessors, bound by the vendor's own privacy policy rather than by any obligation the attorney imposed. So a California lawyer who runs a privileged document through a mainstream enterprise legal AI platform violates the statute as written, no matter how strong the vendor's privacy terms are, because access doesn't run through the attorney's own authorization.

The nondelegation clause has no limiting principle.

The bigger problem is Section 6068.1(a)(2), which states, "An attorney shall not delegate the practice of law to generative artificial intelligence." But the bill's text does not define what the "practice of law" is. Indeed, nothing in the bill tells us whether using AI to run a first pass at research or draft a section counts as delegation — even when the attorney reviews every word, revises the analysis and signs the final product as the sole author.

A lawyer who uses AI as a drafting tool and personally verifies every citation is doing what the profession has always allowed with an associate, a paralegal or a research service. But this bill doesn't say that use is fine. It leaves open whether any AI-assisted drafting or research is the delegation the statute forbids, no matter how much lawyering happens before or after the AI's output hits the page.

Critically, the bill's cost lands first on the people least able to absorb it — solo and small-firm lawyers, legal services organizations, and contingency fee plaintiffs counsel, for whom AI has collapsed the price of first-pass research and drafting. It lands next on California lawyers bidding for national work against out-of-state counsel under no such constraint.

And because the penalty for misinterpreting the bill's scope could potentially be sanctions from the court, most firms will not litigate the statute's meaning. They will simply ban the tools, or at least vastly underuse them. This will thus create a chilling effect on AI usage in the law — thus leaving on the table the substantial benefits and efficiencies AI could otherwise offer.

Arbitrators face an explicit, harder line.

For arbitrators, the ambiguity disappears into something stricter. The bill provides that an "arbitrator shall not delegate any part of their decisionmaking process to any generative artificial intelligence tool." It also states that arbitrators "shall not rely on information generated by generative artificial intelligence outside the record without making appropriate disclosures to the parties beforehand and, as far as practical, allowing the parties to comment on its use." There's no verification safe harbor here.

An arbitrator who uses AI to help review voluminous evidence, organize exhibits or flag inconsistencies across testimony risks violating a statute that bars delegating any part of the process — language broad enough to reach organizational tools that never touch the actual decision.

The problem with this is that most California arbitrators will simply decide to forego using AI, which negatively affects the parties by increasing costs and depriving arbitrators of valuable tools that can make their legal and factual research faster and better.

Practitioners should assume the broad reading and document everything.

The provisions tracking existing duties of competence and confidentiality don't need a statute. Courts were already sanctioning unverified AI citations under rules that predate this bill. Codifying a personal verification duty gives that practice a statutory name. It doesn't change what a lawyer must do.

The provisions going beyond those duties do so without defining their own limits. And these provisions stigmatize AI in a way that will discourage even proper usage.

Firms should thus build supervision and documentation into every AI workflow. And because the governor's inaction is enough to put this bill on the books, the question worth watching is not whether he signs it, but whether other states copy it. Lawyers elsewhere should resist that.

Joshua Wurtzel is a partner at Schlam Stone & Dolan LLP.

The opinions expressed are those of the author(s) and do not necessarily reflect the views of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.